UNITED STATES PATENT APPLICATION

for

MULTI-CARRIER COMMUNICATIONS WITH ADAPTIVE CLUSTER CONFIGURATION AND SWITCHING

Inventors:

Xiaodong Li

Hui Liu

Wenzhong Zhang

Kemin Li

prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

12400 Wilshire Boulevard

Los Angeles, CA 90025-1026

(408) 720-8598

File No.: 005158.P007X

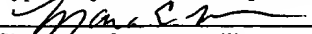# MULTI-CARRIER COMMUNICATIONS WITH ADAPTIVE CLUSTER CONFIGURATION AND SWITCHING

[0001] This patent application is a Continuation-in-part (CIP) of patent application serial no. 09/738,086 filed December 15, 2000, entitled "OFDMA with Adaptive Subcarrier-Cluster Configuration and Selective Loading."

FIELD OF THE INVENTION

[0002] The invention relates to the field of wireless communications; more particularly, the invention relates to multi-cell, multi-subscriber wireless systems using orthogonal frequency division multiplexing (OFDM).

BACKGROUND OF THE INVENTION

[0003] Orthogonal frequency division multiplexing (OFDM) is an efficient modulation scheme for signal transmission over frequency-selective channels. In OFDM, a wide bandwidth is divided into multiple narrow-band subcarriers, which are arranged to be orthogonal with each other. The signals modulated on the subcarriers are transmitted in parallel. For more information, see Cimini, Jr., "Analysis and Simulation of a Digital Mobile Channel Using Orthogonal Frequency Division Multiplexing," IEEE Trans. Commun., vol. COM-33, no. 7, July 1985, pp. 665-75; Chuang and

Sollenberger, "Beyond 3G: Wideband Wireless Data Access Based on OFDM and Dynamic Packet Assignment," IEEE Communications Magazine, Vol. 38, No. 7, pp. 78–87, July 2000.

[0004] One way to use OFDM to support multiple access for multiple subscribers is through time division multiple access (TDMA), in which each subscriber uses all the subcarriers within its assigned time slots. Orthogonal frequency division multiple access (OFDMA) is another method for multiple access, using the basic format of OFDM. In OFDMA, multiple subscribers simultaneously use different subcarriers, in a fashion similar to frequency division multiple access (FDMA). For more information, see Sari and Karam, "Orthogonal Frequency-Division Multiple Access and its Application to CATV Networks," European Transactions on Telecommunications, Vol. 9 (6), pp. 507-516, Nov./Dec. 1998 and Nogueroles, Bossert, Donder, and Zyablov, "Improved Performance of a Random OFDMA Mobile Communication System,", Proceedings of IEEE VTC'98, pp. 2502 –2506.

[0005] Multipath causes frequency-selective fading. The channel gains are different for different subcarriers. Furthermore, the channels are typically uncorrelated for different subscribers. The subcarriers that are in deep fade

for one subscriber may provide high channel gains for another subscriber. Therefore, it is advantageous in an OFDMA system to adaptively allocate the subcarriers to subscribers so that each subscriber enjoys a high channel gain. For more information, see Wong et al., "Multiuser OFDM with Adaptive Subcarrier, Bit and Power Allocation," IEEE J. Select. Areas Commun., Vol. 17(10), pp. 1747-1758, October 1999.

[0006] Within one cell, the subscribers can be coordinated to have different subcarriers in OFDMA. The signals for different subscribers can be made orthogonal and there is little intracell interference. However, with aggressive frequency reuse plan, e.g., the same spectrum is used for multiple neighboring cells, the problem of intercell interference arises. It is clear that the intercell interference in an OFDMA system is also frequency selective and it is advantageous to adaptively allocate the subcarriers so as to mitigate the effect of intercell interference.

[0007] One approach to subcarrier allocation for OFDMA is a joint optimization operation, not only requiring the activity and channel knowledge of all the subscribers in all the cells, but also requiring frequent rescheduling every time an existing subscribers is dropped off the network or a new subscribers is added onto the network. This is often impractical in

real wireless system, mainly due to the bandwidth cost for updating the

subscriber information and the computation cost for the joint optimization.

SUMMARY OF THE INVENTION

[0008] A method and apparatus for allocating subcarriers in an orthogonal frequency division multiple access (OFDMA) system is described. In one embodiment, the method comprises allocating at least one diversity cluster of subcarriers to a first subscriber and allocating at least one coherence cluster to a second subscriber.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The present invention will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the invention, which, however, should not be taken to limit the invention to the specific embodiments, but are for explanation and understanding only.

[0010] **Figure 1A** illustrates subcarriers and clusters.

[0011] **Figure 1B** is a flow diagram of one embodiment of a process for allocating subcarriers.

[0012] **Figure 2** illustrates time and frequency grid of OFDM symbols, pilots and clusters.

[0013] **Figure 3** illustrates subscriber processing.

[0014] **Figure 4** illustrates one example of Figure 3.

[0015] **Figure 5** illustrates one embodiment of a format for arbitrary cluster feedback.

[0016] **Figure 6** illustrates one embodiment of a partition the clusters into groups.

[0017] **Figure 7** illustrates one embodiment of a feedback format for group-based cluster allocation.

[0018] **Figure 8** illustrates frequency reuse and interference in a multi-cell, multi-sector network.

[0019] **Figure 9** illustrates different cluster formats for coherence clusters and diversity clusters.

[0020] **Figure 10** illustrates diversity clusters with subcarrier hopping.

[0021] **Figure 11** illustrates intelligent switching between diversity clusters and coherence clusters depending on subscribers mobility.

[0022] **Figure 12** illustrates one embodiment of a reconfiguration of cluster classification.

[0023] **Figure 13** illustrates one embodiment of a base station.

DETAILED DESCRIPTION OF THE PRESENT INVENTION

[0024] A method and apparatus for allocating subcarriers in an orthogonal frequency division multiple access (OFDMA) system is described. In one embodiment, the method comprises allocating at least one diversity cluster of subcarriers to a first subscriber and allocating at least one coherence cluster to a second subscriber.

[0025] The techniques disclosed herein are described using OFDMA (clusters) as an example. However, they are not limited to OFDMA-based systems. The techniques apply to multi-carrier systems in general, where, for example, a carrier can be a cluster in OFDMA, a spreading code in CDMA, an antenna beam in SDMA (space-division multiple access), etc. In one embodiment, subcarrier allocation is performed in each cell separately. Within each cell, the allocation for individual subscribers (e.g., mobiles) is also made progressively as each new subscriber is added to the system as opposed to joint allocation for subscribers within each cell in which allocation decisions are made taking into account all subscribers in a cell for each allocation.

[0026] For downlink channels, each subscriber first measures the channel and interference information for all the subcarriers and then selects multiple

subcarriers with good performance (e.g., a high signal-to-interference plus noise ratio (SINR)) and feeds back the information on these candidate subcarriers to the base station. The feedback may comprise channel and interference information (e.g., signal-to-interference-plus-noise-ratio information) on all subcarriers or just a portion of subcarriers. In case of providing information on only a portion of the subcarriers, a subscriber may provide a list of subcarriers ordered starting with those subcarriers which the subscriber desires to use, usually because their performance is good or better than that of other subcarriers.

[0027] Upon receiving the information from the subscriber, the base station further selects the subcarriers among the candidates, utilizing additional information available at the base station, e.g., the traffic load information on each subcarrier, amount of traffic requests queued at the base station for each frequency band, whether frequency bands are overused, and/or how long a subscriber has been waiting to send information. In one embodiment, the subcarrier loading information of neighboring cells can also be exchanged between base stations. The base stations can use this information in subcarrier allocation to reduce inter-cell interference.

[0028] In one embodiment, the selection by the base station of the channels

to allocate, based on the feedback, results in the selection of

coding/modulation rates. Such coding/modulation rates may be specified

by the subscriber when specifying subcarriers that it finds favorable to use.

For example, if the SINR is less than a certain threshold (e.g., 12 dB),

quadrature phase shift keying (QPSK) modulation is used; otherwise, 16

quadrature amplitude modulation (QAM) is used. Then the base station

informs the subscribers about the subcarrier allocation and the

coding/modulation rates to use.

[0029] In one embodiment, the feedback information for downlink

subcarrier allocation is transmitted to the base station through the uplink

access channel, which occurs in a short period every transmission time slot,

e.g., 400 microseconds in every 10-millisecond time slot. In one

embodiment, the access channel occupies the entire frequency bandwidth.

Then the base station can collect the uplink SINR of each subcarrier directly

from the access channel. The SINR as well as the traffic load information on

the uplink subcarriers are used for uplink subcarrier allocation.

[0030] For either direction, the base station makes the final decision of

subcarrier allocation for each subscriber.

[0031] In the following description, a procedure of selective subcarrier allocation is also disclosed, including methods of channel and interference sensing, methods of information feedback from the subscribers to the base station, and algorithms used by the base station for subcarrier selections.

[0032] In the following description, numerous details are set forth to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form, rather than in detail, in order to avoid obscuring the present invention.

[0033] Some portions of the detailed descriptions which follow are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic

signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0034] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0035] The present invention also relates to apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computer

selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus.

[0036] The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

[0037] A machine-readable medium includes any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer).

For example, a machine-readable medium includes read only memory ("ROM"); random access memory ("RAM"); magnetic disk storage media; optical storage media; flash memory devices; electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.); etc.

Subcarrier Clustering

[0038] The techniques described herein are directed to subcarrier allocation for data traffic channels. In a cellular system, there are typically other channels, pre-allocated for the exchange of control information and other purposes. These channels often include down link and up link control channels, uplink access channels, and time and frequency synchronization channels.

[0039] Figure 1A illustrates multiple subcarriers, such as subcarrier 101, and cluster 102. A cluster, such as cluster 102, is defined as a logical unit that contains at least one physical subcarrier, as shown in Figure 1A. A cluster can contain consecutive or disjoint subcarriers. The mapping between a cluster and its subcarriers can be fixed or reconfigurable. In the latter case, the base station informs the subscribers when the clusters are redefined. In

one embodiment, the frequency spectrum includes 512 subcarriers and each

cluster includes four consecutive subcarriers, thereby resulting in 128

clusters.

An Exemplary Subcarrier/Cluster Allocation Procedure

[0040] Figure 1B is a flow diagram of one embodiment of a process for

allocation clusters to subscribers. The process is performed by processing

logic that may comprise hardware (e.g., dedicated logic, circuitry, etc.),

software (such as that which runs on, for example, a general purpose

computer system or dedicated machine), or a combination of both.

[0041] Referring to Figure 1B, each base station periodically broadcasts pilot

OFDM symbols to every subscriber within its cell (or sector) (processing

block 101). The pilot symbols, often referred to as a sounding sequence or

signal, are known to both the base station and the subscribers. In one

embodiment, each pilot symbol covers the entire OFDM frequency

bandwidth. The pilot symbols may be different for different cells (or

sectors). The pilot symbols can serve multiple purposes: time and frequency

synchronization, channel estimation and signal-to-interference/noise (SINR)

ratio measurement for cluster allocation.

[0042] Next, each subscriber continuously monitors the reception of the pilot

symbols and measures the SINR and/or other parameters, including inter-

cell interference and intra-cell traffic, of each cluster (processing block 102).

Based on this information, each subscriber selects one or more clusters with

good performance (e.g., high SINR and low traffic loading) relative to each

other and feeds back the information on these candidate clusters to the base

station through predefined uplink access channels (processing block 103).

For example, SINR values higher than 10 dB may indicate good

performance. Likewise, a cluster utilization factor less than 50% may be

indicative of good performance. Each subscriber selects the clusters with

relatively better performance than others. The selection results in each

subscriber selecting clusters they would prefer to use based on the measured

parameters.

[0043] In one embodiment, each subscriber measures the SINR of each

subcarrier cluster and reports these SINR measurements to their base station

through an access channel. The SINR value may comprise the average of the

SINR values of each of the subcarriers in the cluster. Alternatively, the SINR

value for the cluster may be the worst SINR among the SINR values of the

subcarriers in the cluster. In still another embodiment, a weighted

averaging of SINR values of the subcarriers in the cluster is used to generate

an SINR value for the cluster. This may be particularly useful in diversity

clusters where the weighting applied to the subcarriers may be different.

[0044] The feedback of information from each subscriber to the base station

contains a SINR value for each cluster and also indicates the

coding/modulation rate that the subscriber desires to use. No cluster index

is needed to indicate which SINR value in the feedback corresponds to

which cluster as long as the order of information in the feedback is known to

the base station. In an alternative embodiment, the information in the

feedback is ordered according to which clusters have the best performance

relative to each other for the subscriber. In such a case, an index is needed to

indicate to which cluster the accompanying SINR value corresponds.

[0045] Upon receiving the feedback from a subscriber, the base station

further selects one or more clusters for the subscriber among the candidates

(processing block 104). The base station may utilize additional information

available at the base station, e.g., the traffic load information on each

subcarrier, amount of traffic requests queued at the base station for each

frequency band, whether frequency bands are overused, and how long a

subscriber has been waiting to send information. The subcarrier loading

information of neighboring cells can also be exchanged between base stations. The base stations can use this information in subcarrier allocation to reduce inter-cell interference.

[0046] After cluster selection, the base station notifies the subscriber about the cluster allocation through a downlink common control channel or through a dedicated downlink traffic channel if the connection to the subscriber has already been established (processing block 105). In one embodiment, the base station also informs the subscriber about the appropriate modulation/coding rates.

[0047] Once the basic communication link is established, each subscriber can continue to send the feedback to the base station using a dedicated traffic channel (e.g., one or more predefined uplink access channels).

[0048] In one embodiment, the base station allocates all the clusters to be used by a subscriber at once. In an alternative embodiment, the base station first allocates multiple clusters, referred to herein as the basic clusters, to establish a data link between the base station and the subscriber. The base station then subsequently allocates more clusters, referred to herein as the auxiliary clusters, to the subscriber to increase the communication bandwidth. Higher priorities can be given to the assignment of basic

clusters and lower priorities may be given to that of auxiliary clusters. For example, the base station first ensures the assignment of the basic clusters to the subscribers and then tries to satisfy further requests on the auxiliary clusters from the subscribers. Alternatively, the base station may assign auxiliary clusters to one or more subscribers before allocating basic clusters to other subscribers. For example, a base station may allocate basic and auxiliary clusters to one subscriber before allocating any clusters to other subscribers. In one embodiment, the base station allocates basic clusters to a new subscriber and then determines if there are any other subscribers requesting clusters. If not, then the base station allocates the auxiliary clusters to that new subscriber.

[0049] From time to time, processing logic performs retraining by repeating the process described above (processing block 106). The retraining may be performed periodically. This retraining compensates for subscriber movement and any changes in interference. In one embodiment, each subscriber reports to the base station its updated selection of clusters and their associated SINRs. Then the base station further performs the reselection and informs the subscriber about the new cluster allocation. Retraining can be initiated by the base station, and in which case, the base

station requests a specific subscriber to report its updated cluster selection.

Retraining can also be initiated by the subscriber when it observes channel

deterioration.

Adaptive Modulation and Coding

[0050] In one embodiment, different modulation and coding rates are used

to support reliable transmission over channels with different SINR. Signal

spreading over multiple subcarriers may also be used to improve the

reliability at very low SINR.

[0051] An example coding/modulation table is given below in Table 1.

Table 1

Scheme	Modulation	Code Rate
0	QPSK,1/8 Spreading	½
1	QPSK,¼ Spreading	½
2	QPSK,½ Spreading	½
3	QPSK	½
4	8PSK	2/3
5	16QAM	3/4
6	64QAM	5/6

[0052] In the example above, 1/8 spreading indicates that one QPSK

modulation symbol is repeated over eight subcarriers. The

repetition/spreading may also be extended to the time domain. For

example, one QPSK symbol can be repeated over four subcarriers of two

OFDM symbols, resulting also 1/8 spreading.

[0053] The coding/modulation rate can be adaptively changed according to

the channel conditions observed at the receiver after the initial cluster

allocation and rate selection.

Pilot Symbols and SINR Measurement

[0054] In one embodiment, each base station transmits pilot symbols

simultaneously, and each pilot symbol occupies the entire OFDM frequency

bandwidth, as shown in Figures 2A-C. Referring to Figure 2A-C, pilot

symbols 201 are shown traversing the entire OFDM frequency bandwidth

for cells A, B and C, respectively. In one embodiment, each of the pilot

symbols have a length or duration of 128 microseconds with a guard time,

the combination of which is approximately 152 microseconds. After each

pilot period, there are a predetermined number of data periods followed by

another set of pilot symbols. In one embodiment, there are four data periods

used to transmit data after each pilot, and each of the data periods is 152

microseconds.

[0055] A subscriber estimates the SINR for each cluster from the pilot symbols. In one embodiment, the subscriber first estimates the channel response, including the amplitude and phase, as if there is no interference or noise. Once the channel is estimated, the subscriber calculates the interference/noise from the received signal.

[0056] The estimated SINR values may be ordered from largest to smallest SINRs and the clusters with large SINR values are selected. In one embodiment, the selected clusters have SINR values that are larger than the minimum SINR which still allows a reliable (albeit low-rate) transmission supported by the system. The number of clusters selected may depend on the feedback bandwidth and the request transmission rate. In one embodiment, the subscriber always tries to send the information about as many clusters as possible from which the base station chooses.

[0057] The estimated SINR values are also used to choose the appropriate coding/modulation rate for each cluster as discussed above. By using an appropriate SINR indexing scheme, an SINR index may also indicate a particular coding and modulation rate that a subscriber desires to use. Note that even for the same subscribers, different clusters can have different modulation/coding rates.

[0058] Pilot symbols serve an additional purpose in determining interference among the cells. Since the pilots of multiple cells are broadcast at the same time, they will interfere with each other (because they occupy the entire frequency band). This collision of pilot symbols may be used to determine the amount of interference as a worst case scenario. Therefore, in one embodiment, the above SINR estimation using this method is conservative in that the measured interference level is the worst-case scenario, assuming that all the interference sources are on. Thus, the structure of pilot symbols is such that it occupies the entire frequency band and causes collisions among different cells for use in detecting the worst case SINR in packet transmission systems.

[0059] During data traffic periods, the subscribers can determine the level of interference again. The data traffic periods are used to estimate the intra-cell traffic as well as the inter-cell interference level. Specifically, the power difference during the pilot and traffic periods may be used to sense the (intra-cell) traffic loading and inter-cell interference to select the desirable clusters.

[0060] The interference level on certain clusters may be lower, because these clusters may be unused in the neighboring cells. For example, in cell A, with

respect to cluster A there is less interference because cluster A is unused in

cell B (while it is used in cell C). Similarly, in cell A, cluster B will experience

lower interference from cell B because cluster B is used in cell B but not in

cell C.

[0061] The modulation/coding rate based on this estimation is robust to

frequent interference changes resulted from bursty packet transmission.

This is because the rate prediction is based on the worst case situation in

which all interference sources are transmitting.

[0062] In one embodiment, a subscriber utilizes the information available

from both the pilot symbol periods and the data traffic periods to analyze

the presence of both the intra-cell traffic load and inter-cell interference. The

goal of the subscriber is to provide an indication to the base station as to

those clusters that the subscriber desires to use. Ideally, the result of the

selection by the subscriber is clusters with high channel gain, low

interference from other cells, and high availability. The subscriber provides

feedback information that includes the results, listing desired clusters in

order or not as described herein.

[0063] Figure 3 illustrates one embodiment of subscriber processing. The

processing is performed by processing logic that may comprise hardware

(e.g., dedicated logic, circuitry, etc.), software (such as that which runs on, for example, a general purpose computer system or dedicated machine), or a combination of both.

[0064] Referring to Figure 3, channel/interference estimation processing block 301 performs channel and interference estimation in pilot periods in response to pilot symbols. Traffic/interference analysis processing block 302 performs traffic and interference analysis in data periods in response to signal information and information from channel/interference estimation block 301.

[0065] Cluster ordering and rate prediction processing block 303 is coupled to outputs of channel/interference estimation processing block 301 and traffic/interference analysis processing block 302 to perform cluster ordering and selection along with rate prediction.

[0066] The output of cluster ordering processing block 303 is input to cluster request processing block 304, which requests clusters and modulation/coding rates. Indications of these selections are sent to the base station. In one embodiment, the SINR on each cluster is reported to the base station through an access channel. The information is used for cluster selection to avoid clusters with heavy intra-cell traffic loading and/or strong

interference from other cells. That is, a new subscriber may not be allocated use of a particular cluster if heavy intra-cell traffic loading already exists with respect to that cluster. Also, clusters may not be allocated if the interference is so strong that the SINR only allows for low-rate transmission or no reliable transmission at all.

[0067] The channel/interference estimation by processing block 301 is well-known in the art by monitoring the interference that is generated due to full-bandwidth pilot symbols being simultaneously broadcast in multiple cells. The interface information is forwarded to processing block 302 which uses the information to solve the following equation:

$$H_i S_i + I_i + n_i = y_i$$

where S_i represents the signal for subcarrier (freq. band) i, I_i is the interference for subcarrier i, n_i is the noise associated with subcarrier i, and y_i is the observation for subcarrier i. In the case of 512 subcarriers, i may range from 0 to 511. The I_i and n_i are not separated and may be considered one quantity. The interference/noise and channel gain H_i are not know. During pilot periods, the signal S_i representing the pilot symbols, and the observation y_i are knowns, thereby allowing determination of the channel

gain H_i for the case where there is no interference or noise. Once this is known, it may be plugged back into the equation to determine the interference/noise during data periods since H_i, S_i and y_i are all known.

[0068] The interference information from processing blocks 301 and 302 are used by the subscriber to select desirable clusters. In one embodiment, using processing block 303, the subscriber orders clusters and also predicts the data rate that would be available using such clusters. The predicted data rate information may be obtained from a look up table with precalculated data rate values. Such a look up table may store the pairs of each SINR and its associated desirable transmission rate. Based on this information, the subscriber selects clusters that it desires to use based on predetermined performance criteria. Using the ordered list of clusters, the subscriber requests the desired clusters along with coding and modulation rates known to the subscriber to achieve desired data rates.

[0069] Figure 4 is one embodiment of an apparatus for the selection of clusters based on power difference. The approach uses information available during both pilot symbol periods and data traffic periods to perform energy detection. The processing of Figure 4 may be implemented in hardware, (e.g., dedicated logic, circuitry, etc.), software (such as is run

on, for example, a general purpose computer system or dedicated machine),

or a combination of both.

[0070] Referring to Figure 4, a subscriber includes SINR estimation

processing block 401 to perform SINR estimation for each cluster in pilot

periods, power calculation processing block 402 to perform power

calculations for each cluster in pilot periods, and power calculation

processing block 403 to perform power calculations in data periods for each

cluster. Subtractor 404 subtracts the power calculations for data periods

from processing block 403 from those in pilot periods from processing block

402. The output of subtractor 404 is input to power difference ordering (and

group selection) processing block 405 that performs cluster ordering and

selection based on SINR and the power difference between pilot periods and

data periods. Once the clusters have been selected, the subscriber requests

the selected clusters and the coding/modulation rates with processing block

406.

[0071] More specifically, in one embodiment, the signal power of each

cluster during the pilot periods is compared with that during the traffic

periods, according to the following:

$$P_P = P_S + P_I + P_N \, ,$$

$$P_D = \begin{cases} P_N \text{, with no signal and interference} \\ P_S + P_N \text{, with signal only} \\ P_I + P_N \text{, with interference only} \\ P_S + P_I + P_N \text{, with both signal and interference} \end{cases}$$

$$P_P - P_D = \begin{cases} P_S + P_I \text{, with no signal and interference} \\ P_I \text{, with signal only} \\ P_S \text{, with interference only} \\ 0 \text{, with both signal and interference} \end{cases}$$

where P_p is the measured power corresponding to each cluster during pilot periods, P_D is the measured power during the traffic periods, P_S is the signal power, P_I is the interference power, and P_N is the noise power.

[0072] In one embodiment, the subscriber selects clusters with relatively large $P_p / (P_P - P_D)$ (e.g., larger than a threshold such as 10dB) and avoids clusters with low $P_p / (P_P - P_D)$ (e.g., lower than a threshold such as 10dB) when possible.

[0073] Alternatively, the difference may be based on the energy difference between observed samples during the pilot period and during the data traffic period for each of the subcarriers in a cluster such as the following:

$$\Delta_i = \left| y_i^P \right| - \left| y_i^D \right|$$

Thus, the subscriber sums the differences for all subcarriers.

[0074] Depending on the actual implementation, a subscriber may use the following metric, a combined function of both SINR and $P_P - P_D$, to select the clusters:

$$\beta = f(\text{SINR}, P_P / (P_P - P_D))$$

where f is a function of the two inputs. One example of f is weighted averaging (e.g., equal weights). Alternatively, a subscriber selects a cluster based on its SINR and only uses the power difference $P_P - P_D$ to distinguish clusters with similar SINR. The difference may be smaller than a threshold (e.g., 1 dB).

[0075] Both the measurement of SINR and $P_P - P_D$ can be averaged over time to reduce variance and improve accuracy. In one embodiment, a moving-average time window is used that is long enough to average out the statistical abnormity yet short enough to capture the time-varying nature of channel and interference, e.g., 1 millisecond.

Feedback Format for Downlink Cluster Allocation

[0076] In one embodiment, for the downlink, the feedback contains both the indices of selected clusters and their SINR. An exemplary format for arbitrary cluster feedback is shown in Figure 5. Referring to Figure 5, the subscriber provides a cluster index (ID) to indicate the cluster and its associated SINR value. For example, in the feedback, the subscriber provides cluster ID1 (501) and the SINR for the cluster, SINR1 (502), cluster ID2 (503) and the SINR for the cluster, SINR2 (504), and cluster ID3 (505), and the SINR for the cluster, SINR3 (506), etc. The SINR for the cluster may be created using an average of the SINRs of the subcarriers. Thus, multiple arbitrary clusters can be selected as the candidates. As discussed above, the selected clusters can also be ordered in the feedback to indicate priority. In one embodiment, the subscriber may form a priority list of clusters and sends back the SINR information in a descending order of priority.

[0077] Typically, an index to the SINR level, instead of the SINR itself is sufficient to indicate the appropriate coding/modulation for the cluster. For example, a 3-bit field can be used for SINR indexing to indicate 8 different rates of adaptive coding/modulation.

An Exemplary Base Station

[0078] The base station assigns desirable clusters to the subscriber making the request. In one embodiment, the availability of the cluster for allocation to a subscriber depends on the total traffic load on the cluster. Therefore, the base station selects the clusters not only with high SINR, but also with low traffic load.

[0079] Figure 13 is a block diagram of one embodiment of a base station. Referring to Figure 13, cluster allocation and load scheduling controller 1301 (cluster allocator) collects all the necessary information, including the downlink/uplink SINR of clusters specified for each subscriber (e.g., via SINR/rate indices signals 1313 received from OFDM transceiver 1305) and user data, queue fullness/traffic load (e.g., via user data buffer information 1311 from multi-user data buffer 1302). Using this information, controller 1301 makes the decision on cluster allocation and load scheduling for each user, and stores the decision information in a memory (not shown). Controller 1301 informs the subscribers about the decisions through control signal channels (e.g., control signal/cluster allocation 1312 via OFDM transceiver 1305). Controller 1301 updates the decisions during retraining.

[0080] In one embodiment, controller 1301 also performs admission control to user access since it knows the traffic load of the system. This may be performed by controlling user data buffers 1302 using admission control signals 1310.

[0081] The packet data of User 1 ~ N are stored in the user data buffers 1302. For downlink, with the control of controller 1301, multiplexer 1303 loads the user data to cluster data buffers (for Cluster 1 ~ M) waiting to be transmitted. For the uplink, multiplexer 1303 sends the data in the cluster buffers to the corresponding user buffers. Cluster buffer 1304 stores the signal to be transmitted through OFDM transceiver 1305 (for downlink) and the signal received from transceiver 1305. In one embodiment, each user might occupy multiple clusters and each cluster might be shared by multiple users (in a time-division-multiplexing fashion).

Group-Based Cluster Allocation

[0082] In another embodiment, for the downlink, the clusters are partitioned into groups. Each group can include multiple clusters. Figure 6 illustrates an exemplary partitioning. Referring to Figure 6, groups 1-4 are shown with arrows pointing to clusters that are in each group as a result of the

partitioning. In one embodiment, the clusters within each group are spaced

far apart over the entire bandwidth. In one embodiment, the clusters within

each group are spaced apart farther than the channel coherence bandwidth,

i.e. the bandwidth within which the channel response remains roughly the

same. A typical value of coherence bandwidth is 100 kHz for many cellular

systems. This improves frequency diversity within each group and

increases the probability that at least some of the clusters within a group can

provide high SINR. The clusters may be allocated in groups.

[0083] Goals of group-based cluster allocation include reducing the data bits

for cluster indexing, thereby reducing the bandwidth requirements of the

feedback channel (information) and control channel (information) for cluster

allocation. Group-based cluster allocation may also be used to reduce inter-

cell interference.

[0084] After receiving the pilot signal from the base station, a subscriber

sends back the channel information on one or more cluster groups,

simultaneously or sequentially. In one embodiment, only the information

on some of the groups is sent back to the base station. Many criteria can be

used to choose and order the groups, based on the channel information, the

inter-cell interference levels, and the intra-cell traffic load on each cluster.

[0085] In one embodiment, a subscriber first selects the group with the best overall performance and then feedbacks the SINR information for the clusters in that group. The subscriber may order the groups based on their number of clusters for which the SINR is higher than a predefined threshold. By transmitting the SINR of all the clusters in the group sequentially, only the group index, instead of all the cluster indices, needs to be transmitted. Thus, the feedback for each group generally contains two types of information: the group index and the SINR value of each cluster within the group. Figure 7 illustrates an exemplary format for indicating a group-based cluster allocation. Referring to Figure 7, a group ID, ID1, is followed by the SINR values for each of the clusters in the group. This can significantly reduce the feedback overhead.

[0086] Upon receiving the feedback information from the subscriber, the cluster allocator at the base station selects multiple clusters from one or more groups, if available, and then assigns the clusters to the subscriber. This selection may be performed by an allocation in a media access control portion of the base station.

[0087] Furthermore, in a multi-cell environment, groups can have different priorities associated with different cells. In one embodiment, the

subscriber's selection of a group is biased by the group priority, which means that certain subscribers have higher priorities on the usage of some groups than the other subscribers.

[0088] In one embodiment, there is no fixed association between one subscriber and one cluster group; however, in an alternative embodiment there may be such a fixed association. In an implementation having a fixed association between a subscriber and one or more cluster groups, the group index in the feedback information can be omitted, because this information is known to both subscriber and base station by default.

[0089] In another embodiment, the pilot signal sent from the base station to the subscriber also indicates the availability of each cluster, e.g., the pilot signal shows which clusters have already been allocated for other subscribers and which clusters are available for new allocations. For example, the base station can transmit a pilot sequence 1111 1111 on the subcarriers of a cluster to indicate that the cluster is available, and 1111 –1-1-1-1 to indicate the cluster is not available. At the receiver, the subscriber first distinguishes the two sequences using the signal processing methods which are well known in the art, e.g., the correlation methods, and then estimates the channel and interference level.

[0090] With the combination of this information and the channel

characteristics obtained by the subscriber, the subscriber can prioritize the

groups to achieve both high SINR and good load balancing.

[0091] In one embodiment, the subscriber protects the feedback information

by using error correcting codes. In one embodiment, the SINR information

in the feedback is first compressed using source coding techniques, e.g.,

differential encoding, and then encoded by the channel codes.

[0092] Figure 8 shows one embodiment of a frequency reuse pattern for an

exemplary cellular set up. Each cell has hexagonal structure with six sectors

using directional antennas at the base stations. Between the cells, the

frequency reuse factor is one. Within each cell, the frequency reuse factor is

2 where the sectors use two frequencies alternatively. As shown in Figure 8,

each shaded sector uses half of the available OFDMA clusters and each

unshaded sector uses the other half of the clusters. Without loss of

generality, the clusters used by the shaded sectors are referred to herein as

odd clusters and those used by the unshaded sectors are referred to herein

as even clusters.

[0093] Consider the downlink signaling with omni-directional antennas at

the subscribers. From Figure 8, it is clear that for the downlink in the shaded

sectors, Cell A interferes with Cell B, which in turn interferes with Cell C, which in turn interferes with Cell A, namely, A -> B -> C ->A. For the unshaded sectors, Cell A interferes with Cell C, which in turn interferes with Cell B, which in turn interferes with Cell A, namely, A -> C -> B ->A.

[0094] Sector A1 receives interference from Sector C1, but its transmission interferes with Sector B1. Namely, its interference source and the victims with which it interferes are not the same. This might cause a stability problem in a distributed cluster-allocation system using interference avoidance: if a frequency cluster is assigned in Sector B1 but not in Sector C1, the cluster may be assigned in A1 because it may be seen as clean in A1. However, the assignment of this cluster A1 can cause interference problem to the existing assignment in B1.

[0095] In one embodiment, different cluster groups are assigned different priorities for use in different cells to alleviate the aforementioned problem when the traffic load is progressively added to a sector. The priority orders are jointly designed such that a cluster can be selectively assigned to avoid interference from its interference source, while reducing, and potentially minimizing, the probability of causing interference problem to existing assignments in other cells.

[0096] Using the aforementioned example, the odd clusters (used by the shaded sectors) are partitioned into 3 groups: Group 1, 2, 3. The priority orders are listed in Table 2.

Table 2: Priority ordering for the downlink of the shaded sectors.

Priority Ordering	Cell A	Cell B	Cell C
1	Group 1	Group 3	Group 2
2	Group 2	Group 1	Group 3
3	Group 3	Group 2	Group 1

[0097] Consider Sector A1. First, the clusters in Group 1 are selectively assigned. If there are still more subscribers demanding clusters, the clusters in Group 2 are selectively assigned to subscribers, depending on the measured SINR (avoiding the clusters receiving strong interference from Sector C1). Note that the newly assigned clusters from Group 2 to Sector A1 shall not cause interference problem in Sector B1, unless the load in Sector B1 is so heavy that the clusters in both Group 3 and 1 are used up and the clusters in Group 2 are also used. Table 3 shows the cluster usage when less than 2/3 of all the available clusters are used in Sector A1, B1, and C1.

Table 3: Cluster usage for the downlink of the shaded sectors with less than 2/3 of the full load.

Cluster Usage	Cell A	Cell B	Cell C
1	Group 1	Group 3	Group 2
2	Group 2	Group 1	Group 3
3			

[0098] Table 4 shows the priority orders for the unshaded sectors, which are different from those for the shaded sectors, since the interfering relationship is reversed.

Table 4: Priority ordering for the downlink of the unshaded sectors.

Priority Ordering	Cell A	Cell B	Cell C
1	Group 1	Group 2	Group 3
2	Group 2	Group 3	Group 1
3	Group 3	Group 1	Group 2

Intelligent Switching between Coherence and Diversity Clusters

[0099] In one embodiment, there are two categories of clusters: coherence clusters, containing multiple subcarriers close to each other and diversity clusters, containing multiple subcarriers with at least some of the subcarriers spread far apart over the spectrum. The closeness of the multiple subcarriers in coherence clusters is preferably within the channel coherence bandwidth, i.e. the bandwidth within which the channel response remains roughly the same, which is typically within 100 kHz for many cellular systems. On the other hand, the spread of subcarriers in diversity clusters is

preferably larger than the channel coherence bandwidth, typically within

100 kHz for many cellular systems. Of course, the larger the spread, the

better the diversity. Therefore, a general goal in such cases is to maximize

the spread.

[0100] Figure 9 illustrates exemplary cluster formats for coherence clusters

and diversity clusters for Cells A-C. Referring to Figure 9, for cells A-C, the

labeling of frequencies (subcarriers) indicates whether the frequencies are

part of coherence or diversity clusters. For example, those frequencies

labeled 1-8 are diversity clusters and those labeled 9-16 are coherence

clusters. For example, all frequencies labeled 1 in a cell are part of one

diversity cluster, all frequencies labeled 2 in a cell are part of another

diversity cluster, etc., while the group of frequencies labeled 9 are one

coherence cluster, the group of frequencies labeled 10 are another coherence

cluster, etc. The diversity clusters can be configured differently for different

cells to reduce the effect of inter-cell interference through interference

averaging.

[0101] Figure 9 shows example cluster configurations for three neighboring

cells. The interference from a particular cluster in one cell are distributed to

many clusters in other cells, e.g., the interference from Cluster 1 in Cell A are

distributed to Cluster 1, 8, 7, 6 in Cell B. This significantly reduces the interference power to any particular cluster in Cell B. Likewise, the interference to any particular cluster in one cell comes from many different clusters in other cells. Since not all cluster are strong interferers, diversity clusters, with channel coding across its subcarriers, provide interference diversity gain. Therefore, it is advantageous to assign diversity clusters to subscribers that are close (e.g., within the coherent bandwidth) to the cell boundaries and are more subject to inter-cell interference.

[0102] Since the subcarriers in a coherence cluster are consecutive or close (e.g., within the coherent bandwidth) to each other, they are likely within the coherent bandwidth of the channel fading. Therefore, the channel gain of a coherence cluster can vary significantly and cluster selection can greatly improve the performance. On the other hand, the average channel gain of a diversity cluster has less of a degree of variation due to the inherent frequency diversity among the multiple subcarriers spread over the spectrum. With channel coding across the subcarriers within the cluster, diversity clusters are more robust to cluster mis-selection (by the nature of diversification itself), while yielding possibly less gain from cluster selection. Channel coding across the subcarriers means that each codeword contains

bits transmitted from multiple subcarriers, and more specifically, the difference bits between codewords (error vector) are distributed among multiple subcarriers.

[0103] More frequency diversity can be obtained through subcarrier hopping over time in which a subscriber occupies a set of subcarriers at one time slot and another different set of subcarriers at a different time slot. One coding unit (frame) contains multiple such time slots and the transmitted bits are encoded across the entire frame.

[0104] Figure 10 illustrates diversity cluster with subcarrier hopping. Referring to Figure 10, there are four diversity clusters in each of cells A and B shown, with each subcarrier in individual diversity clusters having the same label (1, 2, 3, or 4). There are four separate time slots shown and during each of the time slots, the subcarriers for each of the diversity clusters change. For example, in cell A, subcarrier 1 is part of diversity cluster 1 during time slot 1, is part of diversity cluster 2 during time slot 2, is part of diversity cluster 3 during time slot 3, and is part of diversity cluster 4 during time slot 4. Thus, more interference diversity can be obtained through subcarrier hopping over time, with further interference diversity achieved by using different hopping patterns for different cells, as shown in Figure 10.

[0105] The manner in which the subscriber changes the subcarriers (hopping sequences) can be different for different cells in order to achieve better interference averaging through coding.

[0106] For static subscribers, such as in fixed wireless access, the channels change very little over time. Selective cluster allocation using the coherence clusters achieves good performance. On the other hand, for mobile subscribers, the channel time variance (the variance due to changes in the channel over time) can be very large. A high-gain cluster at one time can be in deep fade at another. Therefore, cluster allocation needs to be updated at a rapid rate, causing significant control overhead. In this case, diversity clusters can be used to provide extra robustness and to alleviate the overhead of frequent cluster reallocation. In one embodiment, cluster allocation is performed faster than the channel changing rate, which is often measured by the channel Doppler rate (in Hz), i.e. how many cycles the channel changes per second where the channel is completely different after one cycle. Note that selective cluster allocation can be performed on both coherence and diversity clusters.

[0107] In one embodiment, for cells containing mixed mobile and fixed subscribers, a channel/interference variation detector can be implemented at

either the subscriber or the base station, or both. Using the detection results,

the subscriber and the base station intelligently selects diversity clusters to

mobile subscribers or fixed subscribers at cell boundaries, and coherence

clusters to fixed subscribers close to the base station. The

channel/interference variation detector measures the channel (SINR)

variation from time to time for each cluster. For example, in one

embodiment, the channel/interference detector measures the power

difference between pilot symbols for each cluster and averages the

difference over a moving window (e.g., 4 time slots). A large difference

indicates that channel/interference changes frequently and subcarrier

allocation may be not reliable. In such a case, diversity clusters are more

desirable for the subscriber.

[0108] Figure 11 is a flow diagram of one embodiment of a process for

intelligent selection between diversity clusters and coherence clusters

depending on subscribers mobility. The process is performed by processing

logic that may comprise hardware (e.g., circuitry, dedicated logic, etc.),

software (such as that which runs on, for example, a general purpose

computer system or dedicated machine), or a combination of both.

[0109] Referring to Figure 11, processing logic in the base station performs channel/interference variation detection (processing block 1101). Processing logic then tests whether the results of the channel/interference variation detection indicate that the user is mobile or in a fixed position close to the edge of the cell (processing block 1102). If the user is not mobile or is not in a fixed position close to the edge of the cell, processing transitions to processing block 1103 where processing logic in the base station selects coherence clusters; otherwise, processing transitions to processing block 1104 in which processing logic in the base station selects diversity clusters.

[0110] In one embodiment, the base station determines whether a subscriber is mobile or fixed by detecting a rate of change of pilot signals, or the normalized channel variation, and determining that the rate of change is greater than a predetermined threshold. The normalized instantaneous difference between channels may be represented as $\dfrac{|H_i - H_{i+1}|^2}{|H_i|^2}$ where H_i represents the channel and i is the index to represent the individual channels.

[0111] The threshold is system dependent. For example, the rate of change is greater than 10% (although any precentage (e.g., 20%) could be used), then

the base station concludes that the subscriber is mobile. In one embodiment, if the constant period in signaling is not greater than a multiple of the round trip delay (e.g., 5 times the round trip delay), then the base station determines that the subscriber is mobile and allocates diversity clusters; otherwise, the base station allocates coherence clusters.

[0112] The selection can be updated and intelligently switched during retraining.

[0113] The ratio/allocation of the numbers of coherence and diversity clusters in a cell depends on the ratio of the population of mobile and fixed subscribers. When the population changes as the system evolves, the allocation of coherence and diversity clusters can be reconfigured to accommodate the new system needs. Figure 12 illustrates a reconfiguration of cluster classification which can support more mobile subscribers than that in Figure 9.

[0114] Whereas many alterations and modifications of the present invention will no doubt become apparent to a person of ordinary skill in the art after having read the foregoing description, it is to be understood that any particular embodiment shown and described by way of illustration is in no way intended to be considered limiting. Therefore, references to details of

various embodiments are not intended to limit the scope of the claims which

in themselves recite only those features regarded as essential to the

invention.